arXiv:1408.3325v1 [cs.CG] 14 Aug 2014

Simultaneous Drawing of Planar Graphs with Right-Angle Crossings and Few Bends

Michael A. Bekos[1], Thomas C. van Dijk[2], Philipp Kindermann[2], Alexander Wolff[2]

[1] Wilhelm-Schickard-Institut für Informatik, Universität Tübingen, Germany
bekos@informatik.uni-tuebingen.de
[2] Lehrstuhl für Informatik I, Universität Würzburg, Germany.
http://www1.informatik.uni-wuerzburg.de/en/staff

Abstract. Given two planar graphs that are defined on the same set of vertices, *a RAC simultaneous drawing* is one in which each graph is drawn planar, there are no edge overlaps and the crossings between the two graphs form right angles. The geometric version restricts the problem to straight-line drawings. It is known, however, that there exists a wheel and a matching which do not admit a geometric RAC simultaneous drawing.
In order to enlarge the class of graphs that admit RAC simultaneous drawings, we allow bends in the resulting drawings. We prove that two planar graphs always admit a RAC simultaneous drawing with six bends per edge each, in quadratic area. For more restricted classes of planar graphs (i.e., matchings, paths, cycles, outerplanar graphs and subhamiltonian graphs), we manage to significantly reduce the required number of bends per edge, while keeping the area quadratic.

1 Introduction

A simultaneous embedding of two planar graphs embeds each graph in a planar way—using the same vertex positions for both embeddings. Edges of one graph are allowed to intersect edges of the other graph. There are two versions of the problem: In the first version, called *Simultaneous Embedding with Fixed Edges* (SEFE), edges that occur in both graphs must be embedded in the same way in both graphs (and hence, cannot be crossed by any other edge). In the second version, these edges can be drawn differently for each of the graphs. Both versions of the problem have a geometric variant where edges must be drawn using straight-line segments.

Simultaneous embedding problems have been extensively investigated over the last few years, starting with the work of Brass et al. [5] on simultaneous straight-line drawing problems. Bläsius et al. [4] recently surveyed the area. For example, it is possible to decide in linear time whether a pair of graphs admits a SEFE or not, if the common graph is biconnected [1]. When actually drawing these simultaneous embeddings, a natural choice is to use straight-line segments. Only very few graphs can be drawn in this way, however, and some existing results need exponential area. For instance, there exist a tree and a path which cannot be drawn simultaneously with straight-line segments [2], and the algorithm for simultaneously drawing a tree and a matching [7] does not provide a polynomial area bound. For the case of edges with bends, that is, polygonal edges, Erten and Kobourov [9] showed that three bends and quadratic area

suffice for any pair of planar graphs (without fixed edges), and that one bend suffices for pairs of trees. Kammer [10] reduced the number of bends to two for the general planar case. In these results, however, the *crossing angles* can be very small.

We suggest a new approach that overcomes the aforementioned problems. We insist that crossings occur at right angles, thereby "taming" them. We do this while drawing on a grid of size $O(n) \times O(n)$ for n-vertex graphs, and we can still draw any pair of planar graphs simultaneously. We do not consider the problem of fixed edges. In a way, our results give a measure for the geometric complexity of simultaneous embeddability for various pairs of graph classes, some of which can be combined more easily (that is, with fewer bends) and some not as easily (needing more bends).

Brightwell and Scheinermann [6] proved that the problem of simultaneously drawing a (primal) embedded graph and its dual always admits a solution if the input graph is a triconnected planar graph. Erten and Kobourov [8] presented an $O(n)$-time algorithm that computes simultaneous drawings of a triconnected planar graph and its dual on an $O(n^2)$ integer grid, where n is the total number of vertices in the graph and its dual. However, these drawings can have non-right angle crossings.

More formally, in this paper we study the *RAC simultaneous drawing problem* (RACSIM *drawing problem*). Let $G_1 = (V, E_1)$ and $G_2 = (V, E_2)$ be two planar graphs on the same vertex set. We say that G_1 and G_2 admit a RACSIM drawing if we can place the vertices on the plane such that (i) each edge is drawn as a polyline, (ii) each graph is drawn planar, (iii) there are no edge overlaps and (iv) crossings between edges in E_1 and E_2 occur at right angles.

Argyriou et al. [3] introduced and studied the geometric version of RACSIM drawing. In particular, they proved that it is always possible to construct a geometric RACSIM drawing of a cycle and a matching in quadratic area, while there exist a wheel and a cycle which do not admit a geometric RACSIM drawing. The problem that we study was left as an open problem.

Our contribution. First, we look at the most general version of the problem: Two planar graphs. (In a simultaneous drawing, certainly both graphs must—individually—be planar.) We give a linear-time algorithm for this case, which produces a drawing in quadratic area with at most six bends per edge. For 2-page book embeddable graphs

Graph classes			Number of bends	Ref.
Planar	+	Planar	$6+6$	Thm. 1
2-page book embeddable	+	2-page book embeddable	$4+4$	Cor. 1
Outerplanar	+	Outerplanar	$3+3$	Thm. 2
Cycle	+	Cycle	$1+1$	Thm. 3
Caterpillar	+	Cycle	$1+1$	Thm. 4
Four Matchings			$1+1+1+1$	Thm. 5
Tree	+	Matching	$1+0$	Thm. 6
Wheel	+	Matching	$2+0$	Thm. 7
Outerpath	+	Matching	$2+1$	Thm. 8

Table 1. A short summary of our results.

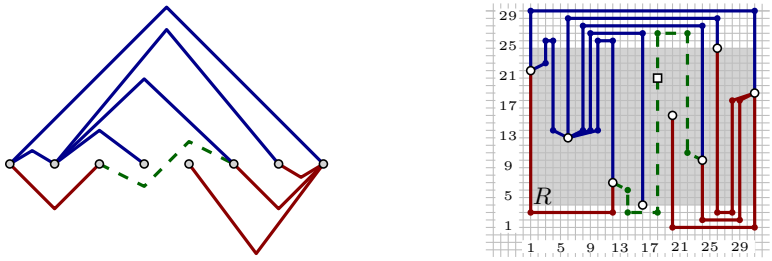

(a) (b)

Fig. 1. (a) A drawing of a planar graph by Kaufmann and Wiese [11] and (b) the drawing by our algorithm with at most six bends per edge. The edge that crosses the spine is drawn dashed. The dummy vertex placed on this edge is drawn as a square.

and outerplanar graphs, we give algorithms that guarantee four and three bends respectively.. Then we turn our attention to graph classes that are more restricted, but for which we can give algorithms that use very few bends. See Table 1 for a full list of results. The main approach in these algorithm is to find linear orders on the vertices of the two graphs and then to compute coordinates for the vertices based on these orders. (See for example also Kaufmann and Wiese [11].)

2 RACSIM Drawings of general graphs

In this section, we study general classes of planar graphs and show how to efficiently construct RACSIM drawings with more than two bends per edge in quadratic area. In particular, we prove two planar graphs on a common set of n vertices admit a RACSIM drawing on an integer grid of size $(14n - 26) \times (14n - 26)$ with six bends per edge (Theorem 1). This result can also be applied to 2-page book embeddable graph, where it gives a RACSIM drawing on an integer grid of size $(11n - 32) \times (11n - 32)$. This also improves the number of bends to four per edge (Corollary 1). If the input is two outerplanar graphs, the algorithm can be improved to get a RACSIM drawing on an integer grid of size $(7n - 10) \times (7n - 10)$, with three bends per edge (Theorem 2).

Theorem 1. *Two planar graphs on a common set of n vertices admit a* RACSIM *drawing on an integer grid of size $(14n - 26) \times (14n - 26)$ with six bends per edge. The drawing can be computed in $O(n)$ time.*

Proof. Let $\mathcal{G}_1 = (V, E_1)$ and $\mathcal{G}_2 = (V, E_2)$ be the two planar graphs. Central in our approach is an algorithm by Kaufmann and Wiese [11], which, given a planar graph $\mathcal{G}$, computes a mapping of the vertices of $\mathcal{G}$ to the points of a point set $\mathcal{P}$ restricted on a horizontal line (called *spine*), such that when $\mathcal{G}$ is embedded on $\mathcal{P}$, each edge of the graph crosses the spine at most once (Note that the algorithm of Kaufmann and Wiese has been used in the past for the simultaneous drawing problem; see [9].) We subdivide all edges of $\mathcal{G}_i$ that cross the spine, by introducing a single *dummy vertex* for each such edge. Let ξ_i be such an embedding of $\mathcal{G}_i$ and denote by E_i^{A} and E_i^{B} the

edges that are drawn completely above and below the spine in ξ_i, respectively. Also, let $\mathcal{G}'_i = (V'_i, E'_i) = (V \cup V_i, E_i^{\mathsf{A}} \cup E_i^{\mathsf{B}})$ be the resulting graph, where V_i contains the dummy vertices. We denote by $\chi_i : V'_i \to \{1, \ldots, |V'_i|\}$ the linear order of the vertices of $\mathcal{G}_i$ along the spine in ξ_i, $i = 1, 2$.

Let $v_1, \ldots, v_{V'_1} \in V'_i$ be the vertices with $\chi_1(v_i) = i$. We place v_1 in the first column. Between two consecutive vertices v_i and v_{i+1}, we reserve several columns for bends of edges incident to v_i and v_{i+1}, in the following order.

(i) One column for the edges $(v_i, \cdot) \in E_2^{\mathsf{A}}$, if one exists.
(ii) One column for every edge $(v_i, v_j) \in E'_1$ with $j > i$.
(iii) One column for every edge $(v_k, v_{i+1}) \in E'_1$ with $k < i + 1$.
(iv) One column for the edges $(v_{i+1}, \cdot) \in E_2^{\mathsf{B}}$, if one exists.

Note that, for (ii) and (iii), we can save some columns because an edge in E_1^{A} and an edge in E_1^{B} can use the same column for their bend. This procedure gives us the x-coordinates of the vertices. Analogously, we can get the y-coordinates of the vertices by rotating the drawing by 90° and applying the procedure to the order χ_2. Let R be the smallest rectangle enclosing all vertices. We draw $\mathcal{G}'_1$ and $\mathcal{G}'_2$ with at most four bends per edge such that all edge segments of $\mathcal{G}'_1$ in R are either vertical or of y-length exactly 1, and all edge segments of $\mathcal{G}'_2$ in R are either horizontal or of x-length exactly 1; see Fig. 1b.

First, we draw the edges $(v_i, v_j) \in E_1^{\mathsf{A}}$ with $i < j$. We draw the edges in a nested order: When we place the edge (v_i, v_j), then there is no edge $(v_k, v_l) \in E_1^{\mathsf{A}}$ with $k \leq i$ and $l \geq j$ that has not already been drawn. Recall that the first column to the right and the first column to the left of every vertex is reserved for the edges in E_1. We draw(v_i, v_j) with at most 4 bends as follows. We start with a slanted segment that has its end point in the row above v_i, and in the first unused column that does not lie to the left of v_i. We follow with a vertical segment to the top that leaves R. We add a horizontal segment above R. In the last unused column that does not lie to the right of v_j, we add a vertical segment that ends one row above v_j. We close the edge with a slanted segment that has its end point in v_j. We draw the edges in $E_1^{\prime B}$ symmetrically with the horizontal segment below R.

Note that this algorithm always uses the top and the bottom port of a vertex v, if there is at least one edge $(v, \cdot)$ in E_1^{A} and E_1^{B}, respectively. Every dummy vertex t has exactly one edge $(t, \cdot)$ in E_1^{A} and E_1^{B}, respectively. Thus, the edges incident to t only use the top and the bottom port. We create a drawing of $\mathcal{G}_1$ with at most 6 bends per edge by removing the dummy vertices from the drawing. In the same way, we create a drawing of $\mathcal{G}_2$ with at most 6 bends per edge.

We will now show that the drawing obtained by combining the drawing of $\mathcal{G}_1$ and $\mathcal{G}_2$ yields a RACSIM drawing. By construction, all segments of E_1 inside R are either vertical segments or slanted segments of x-length at least 2 and y-length exactly 1. All segments of E_2 inside R are either horizontal segments or slanted segments of x-length exactly 1 and y-length at least 2. Thus, the slanted segments can not overlap. Further, all crossings inside R occur between a horizontal and a vertical segment, and thus form right angles. Also, there are no segments in E_1 that lie to the left or to the right of R, and there are no segments in E_2 that lie above or below R. Hence, there are no crossings outside of R and the drawing is a RACSIM drawing.

We will now count the columns used by the drawing. For every vertex in V except the left-most and the right-most, we reserve two additional columns for the edges in E_2; for the remaining two, we only have to reserve one additional column. For every edge in E_1, we need up to 3 columns: One for each end point of the slanted segment at each vertex, and one for the vertical segment that crosses the spine, if it exists. Note that at least one edge per vertex does not need a slanted segment. For every edge in E_2, we need up to 1 column for the vertical segment to the side of R. Since there are at most $3n-6$ edges, our drawing needs $3n-2+3(3n-6)-n+3n-6 = 14n-26$ columns. Analogously, we can show that the algorithm needs $14n-26$ rows, and thus draws the graphs on a $(14n-26) \times (14n-26)$-grid.

Since the algorithm of Kaufmann and Wiese runs in $O(n)$ time, our algorithm also runs in $O(n)$ total time. □

We can improve the results of Theorem 1 for 2-page book embeddable graphs. In a 2-page book embedding, there are no edges that cross the spine. Since these edges are the only ones that need six bends, we can reduce the number of bends per edge to four. Further, the number of columns and rows are reduced by one per edge. This yields the following corollary.

Corollary 1. *Two 2-page book embeddable graphs on a common set of n vertices admit a* RACSIM *drawing on an integer grid of size $(11n-32) \times (11n-32)$ with four bends per edge.*

Theorem 2. *Two outerplanar graphs on a common set of n vertices admit a* RACSIM *drawing on an integer grid of size $(7n-10) \times (7n-10)$ with three bends per edge.*

Proof. It follows by Nash-Williams' [12] formula that every outerplanar graph has arboricity 2, that is, it can be decomposed into two forests. Let $\mathcal{O}_1 = (V, E_1)$ and $\mathcal{O}_2 = (V, E_2)$ be two outerplanar graphs. We draw both graphs as a 2-page book embedding with each forest on one page.

We create a 1-page book embedding for $\mathcal{O}_1$ and $\mathcal{O}_2$. This gives us the order on the x-coordinates and on the y-coordinates, respectively. It follows by Corollary 1 that, by using the algorithm described in the proof of Theorem 1, we create a RACSIM drawing with at most two bends per edge. We will now show how to adjust the algorithm to reduce the number of bends by one.

Let E_1^A and E_1^B be the two forests $\mathcal{O}_1$ is decomposed into. We will draw the edges of E_1^A above the spine and the edges E_1^B below the spine. By rooting the tree in E_1^A, we can direct each edge such that every vertex has exactly one incoming edge. Recall that, in the drawing produced in Theorem 1, one edge per vertex can use the top port. We adjust the algorithm such that every directed edge (v, w) enters the vertex w from the top port. Thus, we draw the edge as follows. We start with a slanted segment of y-length exactly 1. We follow with a vertical segment to the top. We proceed with a horizontal segment that ends directly above w and finish the edge with a vertical segment that enters w from the top port. We use the same approach for the edges in E_1^B. The second outerplanar graph $\mathcal{O}_2$ is drawn analogously.

Since every port of a vertex is only used once, the drawing has no overlaps. We now analyze the number of columns used. For every vertex but the left-most and right-most,

we again reserve two additional columns for the edges in E_2; for the remaining two vertices, we reserve one additional column. However, the edges in E_1 now only need one column for the bend of the single slanted segment. For every edge in E_2, we need up to 1 column for the vertical segment to the side of R. Since there are at most $2n-4$ edges, our drawing needs $3n-2+2n-4+2n-4 = 7n-10$ columns. Analogously, we can show that the algorithm needs $7n-10$ rows. □

3 RACSIM Drawings with one bend per edge

In this section, we study simple classes of planar graphs and show how to efficiently construct RACSIM drawings with one bend per edge in quadratic area. In particular, we prove that two cycles or four matchings (i.e., two classes of graphs of exactly the same size) on a common set of n vertices admit a RACSIM drawing on an integer grid of size $2n \times 2n$; see Theorems 3 and 5, respectively. If the input to our problem is a caterpillar and a cycle, then a RACSIM drawing with one bend per edge is also possible on an integer grid of size $(2n-1) \times 2n$; see Theorem 4. For a tree and a cycle, we can construct a RACSIM drawing with one bend per tree-edge, and no bends in the edges of the matching on an integer grid of size $n \times (n-1)$; see Theorem 6.

Lemma 1. *Two paths on a common set of n vertices admit a* RACSIM *drawing on an integer grid of size $2n \times 2n$ with at most one bend per edge. The drawing can be computed in $O(n)$ time.*

Proof. Let $\mathcal{P}_1 = (V, E_1)$ and $\mathcal{P}_2 = (V, E_2)$ be the two input paths. Following standard practices from the literature (see e.g., Brass et al. [5]), we draw $\mathcal{P}_1$ x-monotone and $\mathcal{P}_2$ y-monotone. This ensures that each of the two paths will eventually be drawn planar. Next, we describe how to compute the exact coordinates of the vertices (and how to route the edges) of $\mathcal{P}_1$ and $\mathcal{P}_2$, such that all (potential) crossings are at right-angles and more importantly there are no edge-segments overlaps.

More precisely, we denote by $p_i : V \to \{1, 2, \ldots, n\}$ the function which maps a vertex of path $\mathcal{P}_i$ to its position in $\mathcal{P}_i$, $i = 1, 2$. Then, a vertex $v \in V$ is drawn at point $(2p_1(v) - 1, 2p_2(v) - 1)$; see Fig. 2a. An edge $(v, v') \in E_1$, with v to the left of v', has a single bend either at point $(x(v), y(v') - 1)$, if $y(v) < y(v')$ or at point $(x(v), y(v') + 1)$, otherwise. On the other hand, an edge $(v, v') \in E_2$, with v above v', also has a single bend either at point $(x(v) + 1, y(v'))$, if $x(v) < x(v')$ or at point $(x(v) - 1, y(v'))$, otherwise.

Clearly, the area required by the drawing is $(2n-1) \times (2n-1)$. From left to right, the edges of $\mathcal{P}_1$ leave the vertices vertically and enter them "diagonally". Similarly, from bottom to top, the edges of $\mathcal{P}_2$ leave the vertices horizontally and enter them "diagonally". In addition, a non-rectilinear edge-segment is drawn between two consecutive (horizontal or vertical) grid lines. Hence, it cannot be involved in crossings or overlaps. Since $\mathcal{P}_1$ and $\mathcal{P}_2$ are x- and y-monotone, respectively, it follows that all (potential) crossings must involve a vertical edge-segment of $\mathcal{P}_1$ and a horizontal edge-segment of $\mathcal{P}_2$, which clearly yields right-angles at the crossing points. □

We say that an edge uses the bottom (left/right/top, resp.) *port* of a vertex if it enters the vertex from the bottom (left/right/top, resp.).

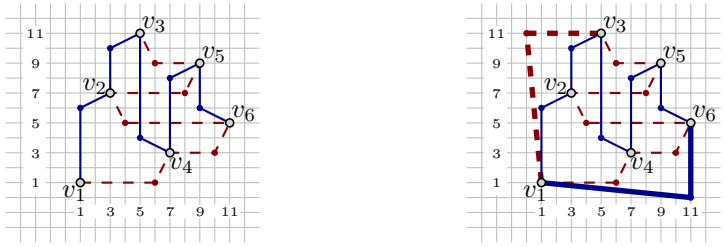

(a) Two paths: $\mathcal{P}_1$ (solid) and $\mathcal{P}_2$ (dashed) (b) Two cycles: $\mathcal{C}_1$ (solid) and $\mathcal{C}_2$(dashed)

Fig. 2. RACSIM drawings with one bend per edge of: (a) two paths and (b) two cycles.

Theorem 3. *Two cycles on a common set of n vertices admit a* RACSIM *drawing on an integer grid of size $2n \times 2n$ with at most one bend per edge. The drawing can be computed in $O(n)$ time.*

Proof. Let $\mathcal{C}_1 = (V, E_1)$ and $\mathcal{C}_2 = (V, E_2)$ be the two input cycles and let $v \in V$ be an arbitrary vertex. We temporarily delete one edge from each of the two cycles incident to vertex v; say $(v, w_1) \in E_1$ from $\mathcal{C}_1$ and $(v, w_2) \in E_2$ from $\mathcal{C}_2$ (refer the bold-drawn edges of Figure 2b). This results into two paths, say $\mathcal{P}_1$ and $\mathcal{P}_2$, with endpoints v and w_1, and, v and w_2, respectively. We employ the algorithm supporting Lemma 1 to construct a RACSIM drawing of $\mathcal{P}_1$ and $\mathcal{P}_2$ on an integer grid of size $(2n-1) \times (2n-1)$. In the resulting drawing, vertex v is placed at the bottom-left corner of the bounding box containing the drawing, w_1 along its right side, and w_2 along its top side. By construction, the bottom port of vertex w_1 and the left port of vertex w_2 are both unoccupied. Hence, the edges (v, w_1) and (v, w_2) that form $\mathcal{C}_1$ and $\mathcal{C}_2$ can be drawn with a single bend each at points $(2n-1, 0)$ and $(0, 2n-1)$, respectively; see Figure 2b. Clearly, none of them is involved in crossings, while the total area of the drawing gets larger by a single unit at each dimension. □

Theorem 4. *A caterpillar and a cycle on a common set of n vertices admit a* RACSIM *drawing on an integer grid of size $(2n-1) \times 2n$ with at most one bend per edge. The drawing can be computed in $O(n)$ time.*

Proof. We denote by $\mathcal{A} = (V, E_{\mathcal{A}})$ and $\mathcal{C} = (V, E_{\mathcal{C}})$ the caterpillar and the cycle, respectively. Starting from a spine vertex that is endpoint of the spine of $\mathcal{A}$, we perform a BFS traversal on $\mathcal{A}$ assuming that we first visit all leg vertices incident to a spine vertex before visiting its neighboring vertex along the spine. Assume, without loss of generality, that $V = \{v_1, v_2, \ldots, v_n\}$ is the ordered vertex set implied by the aforementioned traversal of $\mathcal{A}$; see Figure 3a. As in the proof of Theorem 3, we temporarily delete an edge of $\mathcal{C}$ (incident to the first vertex of V; refer the bold-drawn edge of Figure 3a) and obtain a path, say $\mathcal{P} = (V, E_{\mathcal{P}})$. Let $p : V \to \{1, 2, \ldots, n\}$ be a function which maps a vertex of path $\mathcal{P}$ to its position in $\mathcal{P}$. For $i = 1, 2, \ldots, n$, we draw vertex v_i at point $(2i-1, 2p(v_i)-1)$. An edge $(v, v') \in E_{\mathcal{P}}$, with v above v', has a single bend either at point $(x(v)+1, y(v'))$, if $x(v) < x(v')$, or at point $(x(v)-1, y(v'))$, otherwise. On the other hand, an edge $(v, v') \in E_{\mathcal{A}}$, with v to the left of v', also has a single bend either at point $(x(v'), y(v)+1)$, if $y(v) < y(v')$ or at point $(x(v'), y(v)-1)$, otherwise.

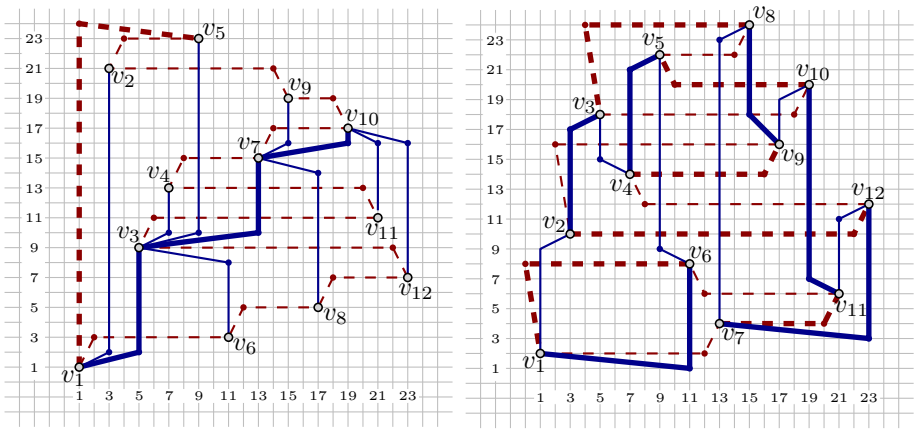

Fig. 3. RACSIM drawings with one bend per edge of: (a) a caterpillar $\mathcal{A}$ (solid; its spine is drawn bold) and a cycle $\mathcal{C}$ (dashed); (b) four matchings $\mathcal{M}_1$ (solid-plain), $\mathcal{M}_2$ (solid-bold), $\mathcal{M}_3$ (dashed-plain) and $\mathcal{M}_4$ (dashed-bold).

The approach described above ensures that $\mathcal{P}$ is drawn y-monotone, hence planar. The spine of $\mathcal{A}$ is drawn x-monotone. The legs of a spine vertex of $\mathcal{A}$ are drawn "to the right" of their parent spine vertex and "to the left" of the next vertex along the spine. Hence, $\mathcal{A}$ is drawn planar as well. The non-rectilinear edge-segments of $\mathcal{A}$ are of y-length one, while the non-rectilinear edge-segments of $\mathcal{P}$ are of x-length one. Thus, they cannot be involved in crossings, which implies that all (potential) crossings form right angles.

Now, it remains to describe how to draw the edge that we removed in order to transform cycle $\mathcal{C}$ to path $\mathcal{P}$. By construction, this edge connects vertex v_1 (which is drawn at the bottom-left corner of the bounding box containing the drawing) with the vertex drawn at the top side of the bounding box containing the drawing. As the top port of v_1 is unoccupied, if this edge bends at $(1, 2n)$, then it is not involved in crossings; see Figure 3a. The total area required by the drawing is $(2n-1) \times 2n$. □

Theorem 5. *Four matchings on a common set of n vertices admit a* RACSIM *drawing on an integer grid of size $2n \times 2n$ with at most one bend per edge. The drawing can be computed in $O(n)$ time.*

Proof. Let $\mathcal{M}_1 = (V, E_1)$, $\mathcal{M}_2 = (V, E_2)$, $\mathcal{M}_3 = (V, E_3)$ and $\mathcal{M}_4 = (V, E_4)$ be the input matchings. Let also $\mathcal{M}_{1,2} = (V, E_1 \cup E_2)$ and $\mathcal{M}_{3,4} = (V, E_3 \cup E_4)$. Since $\mathcal{M}_1$ and $\mathcal{M}_2$ are defined on the same vertex set, $\mathcal{M}_{1,2}$ is a 2-regular graph. Thus, each connected component of $\mathcal{M}_{1,2}$ corresponds to a cycle of even length which alternates between edges of $\mathcal{M}_1$ and $\mathcal{M}_2$; see Figure 3b. The same holds for $\mathcal{M}_{3,4}$. W.l.o.g. we further assume that $\mathcal{M}_{1,2} \cup \mathcal{M}_{3,4}$ is a connected graph. We seek to draw $\mathcal{M}_{1,2}$ x-monotone and $\mathcal{M}_{3,4}$ y-monotone.

We start with choosing an arbitrary vertex $v \in V$. Let $\mathcal{C}$ be the cycle of $\mathcal{M}_{1,2}$ containing vertex v. We determine the x-coordinates of the vertices of $\mathcal{C}$ by traversing

cycle $\mathcal{C}$ in some particular direction (starting from vertex v) and assigning to each vertex of $\mathcal{C}$ its discovery time as the x-coordinate. Next, we determine the y-coordinates of the vertices of all cycles, say $\mathcal{C}_1, \mathcal{C}_2, \ldots \mathcal{C}_k$, of $\mathcal{M}_{3,4}$ that have at least one vertex with a determined x-coordinate. We order these cycles as follows: For a pair of cycles $\mathcal{C}_i$ and $\mathcal{C}_j$, $\mathcal{C}_i$ precedes $\mathcal{C}_j$ if and only if $\mathcal{C}_i$ contains a vertex (which we call *anchor* vertex of $\mathcal{C}_i$) with a determined x-coordinate strictly smaller that the corresponding ones of all vertices of $\mathcal{C}_j$, $i, j \in \{1, 2, \ldots, k\}$. Assume w.l.o.g. that $\mathcal{C}_1 \rightarrow \mathcal{C}_2 \rightarrow \ldots \rightarrow \mathcal{C}_k$ is the computed order. In what follows, we start with the first cycle $\mathcal{C}_1$ of the computed order and determine the y-coordinates of its vertices. To do so, we traverse $\mathcal{C}_1$ in some particular direction (starting from its anchor vertex) and assign to each vertex of $\mathcal{C}_1$ the discovery time as y-coordinate. We proceed similarly with the remaining cycles, assuming that $\mathcal{C}_{i+1}$ is placed directly above $\mathcal{C}_i$, $i = 1, 2, \ldots, k-1$.

Now, there are no vertices with a determined x-coordinate but without a determined y-coordinate. However, the other way around is possible, i.e., there might exist vertices with a determined y-coordinate, but without a determined x-coordinate. If this is the case, we repeat the aforementioned procedure to determine the x-coordinates of the vertices of all cycles of $\mathcal{M}_{1,2} \setminus \mathcal{C}$ that have at least one vertex with a determined y-coordinate. Since at each step of our algorithm the number of vertices that have an undetermined x- or y-coordinate is reduced by the size of the cycles participating in this step and $\mathcal{M}_{1,2} \cup \mathcal{M}_{3,4}$ is a connected graph, our algorithm guarantees that all vertices will be assigned both x and y-coordinate.

Next, we describe how to route the edges of cycles in $\mathcal{M}_{1,2}$ and $\mathcal{M}_{3,4}$, respectively. We draw each edge of $\mathcal{M}_{1,2}$ by using a vertical edge-segment at its left end-vertex, and a slightly slanted horizontal segment at its right end-vertex; see Figure 3b. Note that the slanted edge-segments are all of x-length 1, so they can not be intersected by vertical edge-segments. In addition, the last edge of each cycle (referred to as *closing edge*) is drawn with a vertical edge-segment incident to its right end-point directed downwards below all vertices of the cycle, and a slightly slanted horizontal segment to the its left end-vertex. Similarly, we draw the edges of $\mathcal{M}_{3,4}$; see Figure 3b for an illustration. Our choice of coordinates guarantees that the x-coordinates of the cycles of $\mathcal{M}_{1,2}$ and the y-coordinates of the cycles of $\mathcal{M}_{3,4}$ form disjoint intervals. Thus, the area below a cycle of $\mathcal{M}_{1,2}$ and the area to the left of a cycle of $\mathcal{M}_{3,4}$ are free from vertices. Hence, the slanted segments of the closing edges can not have a crossing that does violate the RAC restriction; see Fig. 3b. □

Theorem 6. *A tree and a matching on a common set of n vertices admit a* RACSIM *drawing on an integer grid of size $n \times (n-1)$ with one bend per tree-edge, and no bends in the edges of the matching. The drawing can be computed in $O(n)$ time.*

Sketch of Proof. We inductively place each matching in one row. In every step, we decide whether to add the next matching to the stack at the top or at the bottom. We determine the x-coordinates of the matching by using a specific post-order. An illustration is given in Fig.4; a detailed proof is given in the appendix. □

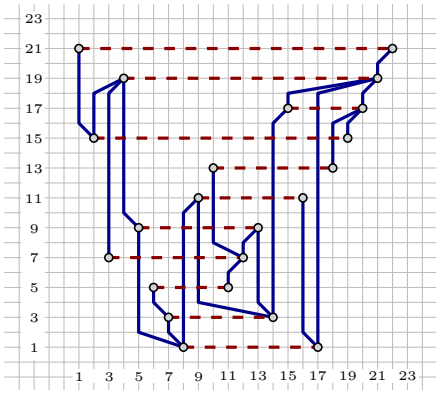

Fig. 4. A RACSIM drawing of a tree (solid) and a matching (dashed)

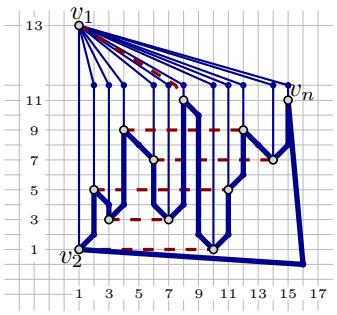

Fig. 5. A RACSIM drawing of a wheel (solid; its rim is drawn bold) and a matching (dashed)

4 RACSIM Drawings with two bends per edge

In this section, we study more complex classes of planar graphs and show how to efficiently construct RACSIM drawings with two bends per edge in quadratic area. In particular, we prove that a wheel and a matching on a common set of n vertices admit a RACSIM drawing on an integer grid of size $(1.5n - 1) \times (n + 2)$ with two bends per edge and no bends, respectively; see Theorem 7. If the input to our problem is an outerpath and a matching, then a RACSIM drawing with two bends per edge and no bends, respectively, is also possible on an integer grid of size $(3n - 2) \times (3n - 2)$; see Theorem 8.

Theorem 7. *A wheel and a matching on a common set of n vertices admit a* RACSIM *drawing on an integer grid of size $(1.5n - 1) \times (n + 2)$ with two bends per edge and no bends, respectively. The drawing can be computed in $O(n)$ time.*

Proof. We denote the wheel by $\mathcal{W}$ and the matching by $\mathcal{M}$. Let the common vertex set of $\mathcal{W}$ and $\mathcal{M}$ be $V = \{v_1, v_2, \ldots, v_n\}$, where $n \geq 4$. If v_1 is the center of $\mathcal{W}$ and $\mathcal{C}_w : v_2 \to v_3 \to \ldots \to v_n \to v_2$ is the rim of $\mathcal{W}$, then $E(\mathcal{W}) = \{(v_i, v_{i+1});\ i = 1, \ldots, n-1\} \cup \{(v_n, v_2)\} \cup \{(v_1, v_i);\ i = 2, \ldots, n\}$. Let $\mathcal{M}'$ be the matching $\mathcal{M}$ without the edge incident to v_1.

We first compute the x-coordinates of the vertices of $\mathcal{W} \cup \mathcal{M}'$, such that $\mathcal{C}_w - \{(v_n, v_2)\}$ is x-monotone (if drawn with straight-line edges). More precisely, for $i = 2, \ldots, n$ we set $x(v_i) = 2i - 3$. The y-coordinates of the vertices of $\mathcal{W} \cup \mathcal{M}'$ are computed based on matching $\mathcal{M}'$. W.l.o.g., we assume that the first edge of $\mathcal{M}'$ is the one incident to v_2; the remaining ones are next in some order. Let $k < n/2$ be the number of matching edges in $\mathcal{M}'$. Then, the end-points of the i-th edge of $\mathcal{M}'$ have y-coordinate $2i - 1$, $i = 1, 2, \ldots, k$. Next, we assign the y-coordinate $2k + 1$ to the vertices incident to the rim without a matching edge. Finally, the center v_1 of $\mathcal{W}$ is located at point $(1, 2k + 3)$.

We now have to describe where each edge of $\mathcal{W}$ bends, as $\mathcal{M}'$ is drawn bendless. We draw the spikes $(v_1, v_i), i = 2, \ldots, n$ of $\mathcal{W}$ with exactly one bend at point $(x(v_i), 2k + 2)$. Since vertex v_2 is left-most in the constructed drawing, we can save the bend of spike (v_0, v_1). We draw the edges (v_i, v_{i+1}) of $\mathcal{C}_w - \{(v_n, v_1)\}$ from left to right. If $y(v_{i+1}) > y(v_i)$, we draw the edge with a single bend at $(x(v_{i+1}), y(v_i) + 1)$. If $y(v_{i-1}) > y(v_i) > y(v_{i+1})$, we draw the edge with a single bend at the point $(x(v_{i+1}), y(v_i) - 1)$. If $y(v_i) > y(v_{i-1}), y(v_{i+1})$, the bottom port at v_i is already used. Thus, we draw the edge with two bends at the point $(x(v_{i+1}), y(v_i) - 1)$ and the point $(x(v_{i+1}), y(v_{i+1}) + 1)$. Finally, we delete the unused columns; see Figure 5.

Our approach ensures that $\mathcal{C}_w - \{(v_n, v_2)\}$ is drawn (non-strictly) x-monotone, hence planar. The last edge (v_n, v_2) of $\mathcal{C}_w$ "surrounds" the drawing; so, it is crossing-free. Further, the spikes are not involved with crossings with the rim of $\mathcal{W}$, as they are drawn in its exterior. So, $\mathcal{W}$ is drawn planar, as desired. On the other hand, all edges of $\mathcal{M}'$ are drawn as horizontal, non-overlapping line-segments. So, $\mathcal{M}'$ is planar as well. The non-rectilinear edge-segments of $\mathcal{W} - (v_n, v_2)$ are of y-length one. So, they cannot be crossed by the edges of $\mathcal{M}'$. As the edge (v_n, v_2) is not involved in crossings, it follows that all (potential) crossings between $\mathcal{W}$ and $\mathcal{M}'$ form right angles.

Finally, we have to insert the matching edge $(v_1, v_i) = \mathcal{M} - \mathcal{M}'$. Since v_i is not incident to a matching edge in $\mathcal{M}'$, it is placed above all matching edges. Then $(v_1, v_i) \in \mathcal{W}$ does not cross a matching edge, so we can use this edge as a double edge.

We will now prove the area bound of the drawing algorithm. First, we count the rows used. Since we remove the matching edge incident to v_1, the matching $\mathcal{M}'$ has $k \leq n/2 - 1 < n/2$ matching edges. We place the bottommost vertex in row 1 and the topmost vertex, that is vertex v_1, in row $2k + 3$. We add one extra bend in row 0 for the edge (v_n, v_2). Thus, our drawing uses $2k + 3 + 1 \leq n + 2$ rows. Next, we count the columns used. The vertices $v_2, \ldots, v_n$ are each placed in their own column. Every spike has exactly one bend in the column of a vertex. An edge (v_i, v_{i+1}) of rim $\mathcal{W}$ has exactly one bend in a vertex column, except for the case where $y(v_i) > y(v_{i-1}), y(v_{i+1})$, in which it needs an extra bend between v_i and v_{i+1}, $i = 1, \ldots, n - 1$. Clearly, there can be at most $n/2 - 1$ vertices satisfying this condition. Since the edge (v_n, v_2) uses an extra column to the right of v_n, our drawing uses $(n - 1) + (n/2 - 1) + 1 = 1.5n - 1$ columns. □

Theorem 8. *An outerpath and a matching on a common set of n vertices admit a* RAC-SIM *drawing on an integer grid of size $(3n - 2) \times (3n - 2)$ with two bends per edge and one bend, respectively. The drawing can be computed in $O(n)$ time.*

Sketch of Proof. Augment the outerpath to maximal outerpath. Removing its outer-cycle, the result is a caterpillar, which determines the x-coordinates of the vertices. The y-coordinates are computed so that the matching is planar. An illustration is given in Fig.6; a detailed proof is given in the appendix. □

5 Conclusions and Open Problems

In this paper, we have studied RAC simultaneous drawings with few bends per edge. We proved that two planar graphs always admit a RAC simultaneous drawing with at most

six bends per edge. For more restricted classes of graphs, we drastically improved the number of bends per edge. All of these drawings are within quadratic area. The results presented in this paper raise several questions that remain open, such as the following.

1. Is it possible to reduce the number of bends per edge for the classes of graphs that we presented in this paper?
2. What additional non-trivial classes of graphs admit a RACSIM drawing with better-than-general number of bends?
3. As a variant of the problem, it might be possible to reduce the required number of bends per edge by relaxing the strict constraint that intersections must be right-angle and instead ask for drawings that have close to optimal crossing resolution.
4. The computational complexity of the general problem remains open: Given two or more planar graphs on the same set of vertices and an integer k, is there a RAC-SIM drawing in which each graph is drawn with at most k bends per edge and the crossing are at right angles?

References

1. Patrizio Angelini, Giuseppe Di Battista, Fabrizio Frati, Maurizio Patrignani, and Ignaz Rutter. Testing the simultaneous embeddability of two graphs whose intersection is a biconnected or a connected graph. *J. Discrete Algorithms*, 14:150–172, 2012.
2. Patrizio Angelini, Markus Geyer, Michael Kaufmann, and Daniel Neuwirth. On a tree and a path with no geometric simultaneous embedding. *J. Graph Algorithms Appl.*, 16(1):37–83, 2012.
3. Evmorfia N. Argyriou, Michael A. Bekos, Michael Kaufmann, and Antonios Symvonis. Geometric RAC simultaneous drawings of graphs. *J. Graph Algorithms Appl.*, 17(1):11–34, 2013.
4. Thomas Bläsius, Stephen G. Kobourov, and Ignaz Rutter. Simultaneous embedding of planar graphs. In Roberto Tamassia, editor, *Handbook of Graph Drawing and Visualization*, chapter 11, pages 349–381. CRC Press, 2013.
5. Peter Brass, Eowyn Cenek, Cristian A. Duncan, Alon Efrat, Cesim Erten, Dan P. Ismailescu, Stephen G. Kobourov, Anna Lubiw, and Joseph S.B. Mitchell. On simultaneous planar graph embeddings. *Comput. Geom. Theory Appl.*, 36(2):117–130, 2007.
6. Graham Brightwell and Edward R. Scheinerman. Representations of planar graphs. *SIAM Journal Discrete Mathematics*, 6(2):214–229, 1993.
7. Sergio Cabello, Marc van Kreveld, Giuseppe Liotta, Henk Meijer, Bettina Speckmann, and Kevin Verbeek. Geometric simultaneous embeddings of a graph and a matching. *J. Graph Algorithms Appl.*, 15(1):79–96, 2011.
8. Cesim Erten and Stephen G. Kobourov. Simultaneous embedding of a planar graph and its dual on the grid. *Theory Computing Systems*, 38(3):313–327, 2005.
9. Cesim Erten and Stephen G. Kobourov. Simultaneous embedding of planar graphs with few bends. *J. Graph Algorithms Appl.*, 9(3):347–364, 2005.
10. Frank Kammer. Simultaneous embedding with two bends per edge in polynomial area. In Lars Arge and Rusins Freivalds, editors, *Proc. SWAT 2006*, volume 4059 of *LNCS*, pages 255–267. Springer-Verlag, 2006.
11. Michael Kaufmann and Roland Wiese. Embedding vertices at points: Few bends suffice for planar graphs. *J. Graph Algorithms Appl.*, 6(1):115–129, 2002.
12. Crispin St. John Alvah Nash-Williams. Decomposition of finite graphs into forests. *Journal of the London Mathematical Society*, 1(1):12–12, 1964.

Appendix

Theorem 6. *A tree and a matching on a common set of n vertices admit a* RACSIM *drawing on an integer grid of size $n \times (n-1)$ with one bend per tree edge and no bends in the edges of the matching. The drawing can be computed in $O(n)$ time.*

Proof. We use an algorithm inspired by the algorithm for drawing a geometric simultaneous embedding of a tree and a matching by Cabello et al. [7]. See Fig. 4 for an example drawing. First of all, we root the tree at an arbitrary leaf, getting a directed tree. We will get the x-coordinates of the vertices from a particular post-order of this directed tree, that is, all children of a vertex are placed to its left. All subtrees of a vertex will be drawn in disjoint x-intervals. We assume that the matching is perfect, otherwise we add dummy edges to make it perfect that are removed until the drawing algorithm is finished.

Each edge of the matching is drawn horizontally, at an odd y-coordinate between 1 and n, inclusive. The ordering of the edges is constructed iteratively as follows. In every step, the algorithm picks an edge and puts it either into the *top group*, or into the *bottom group*. Within the top group, the edges are assigned from top to bottom; in the bottom group, the edges are assigned from bottom to top. Finally, the top group is placed above the bottom group. Once this procedure has picked an edge, the edge and its vertices are called *placed*.

Our algorithm inductively places the matching edges either in the top or in the bottom group. It starts by putting the matching edge containing the root of the tree in the top group and then proceeds as follows. We partition the edges of the tree into a set of subtrees by cutting at the placed vertices. If there exists a subtree with three placed vertices, call the vertex that lies on the three paths between these placed vertices a *splitter*. If the algorithm encounters a splitter, it selects the matching edge incident to the splitter (that edge is unplaced, since a splitter is necessarily unplaced), and places it next. Otherwise, it finds an unplaced vertex that has a tree edge to a placed vertex, and adds its matching edge to the top group. This creates at most one splitter, since one of the placed vertices is adjacent to a vertex that was already placed.

It has been shown by Cabello et al. that placing a vertex in this way creates at most one splitter, and that placing a splitter does not create a new one. In every step of the algorithm, we place two vertices. The first one is either a splitter, or adjacent to a placed vertex; thus, it can not create a new splitter. Hence, when we place a new matching edge, there will be at most one splitter in the graph.

We now describe how to determine whether to add a matching edge that contains a splitter to the top, or the bottom group. Assume that the vertex v is a splitter. We start with drawing the root, so every vertex lies in a subtree that has its local root placed. Let $T_1(v), \ldots, T_k(v)$ be the subtrees rooted in the children of v. Then, v is a splitter if and only if there are two subtrees $T_i(v)$ and $T_j(v)$ with at least one placed vertex, one of which has been placed in the last step. W.l.o.g., let $w \in T_i(v)$ be this vertex. If w was added to the top group, we add v to the bottom group; otherwise, we add it to the top group. Thus, the vertices of all trees $T_l(v)$ with $l \neq i$ will be placed on the same side (with respect to the y-coordinate) of v.

Now, we describe how to get the x-coordinate of the vertices. To that end, we will compute an order for the subtrees of every vertex. This order determines the disjoint intervals on the x-coordinates that the vertices in this subtree will use. Let v be a vertex that is placed in an induction step. We traverse the path from v to the root of the tree up to (and including) the first placed vertex. For every vertex w on this path, we determine the order of the subtrees of its children as follows. Let $T_1(w), \ldots, T_k(w)$ be the subtrees rooted in the children of w, with $v \in T_1(w)$. If v is the only placed successor of vertex w, then we assign the order $x(T_1(w)) < \ldots < x(T_k(w))$ on the x-coordinates of the subtrees; otherwise, an order has already been determined. When the algorithm is done, we know the order of the subtrees for every vertex

We draw the edges of the tree with one horizontal-slanted segment, that is, a segment of y-length exactly 1, at the parent, and a vertical segment at the child. Since we draw the edges of the matching horizontally without bends, and since the slanted segments are drawn between the horizontal grid lines, there can only be crossings between vertical segments of the tree and horizontal segments of the matching. Thus, all crossings between the tree and the matching are at a right angle.

It remains to show that the drawing of the tree is itself plane. Since tree consists of vertical segments and slanted segments, obviously any intersection has to be between a slanted and a vertical segment. We will argue that if the tree self-intersects, then it was not constructed according to the above algorithm. Let v be a vertex of the tree. We will show that the outgoing edges of v do not induce a crossing. Since the subtree rooted in v is assigned an x-interval that contains no other vertices, these edges can only cross other edges in this subtree. Consider the step of the algorithm in which v is placed.

First, assume that v was is not a splitter. If no successor of v has been placed so far, then all successors of v will be placed on the same side of v and therefore not induce a crossing. Otherwise, all placed successors of v are located in the same subtree $T_1(v)$ rooted in a child of v. By construction, $T_1(v)$ is placed to the left of the other subtrees rooted in a child of v. Thus, no edge incident to v is drawn inside the x-interval assigned to $T_i(v)$. The vertices in the other subtrees are yet to be placed, so they will all be on the same side of v and therefore not induce a crossing.

Now, assume that v is a splitter. Then there is a vertex w that was placed in the previous and lies in the subtree of v. Further, there is one subtree $T_1(v)$ rooted in a child of v in which at least one vertex has been placed before w. By construction, $T_1(v)$ is placed to the left of the other subtrees rooted in a child of v. Recall that v is placed in the group opposite of w. Thus, w and all unplaced successors of v lie on the same side of v and therefore will not induce a crossing. This concludes the proof.

Finally, we will prove the area and running time bounds. We use two rows per matching edge, except the top matching edge that only needs one row. Thus, our drawing needs $(n-1)$ rows. In every column, we place exactly one vertex, so the drawing needs n columns. As for the running time, the algorithm to place the vertices clearly requires only constant running time per vertex, with the exception of traversing the tree upwards to determine the x-order of the subtrees. For that, however, we traverse every edge only once, since we stop at the first placed vertex. Thus, the running time of the algorithm is $O(n)$. □

Theorem 8. *An outerpath and a matching on a common set of n vertices admit a* RAC-SIM *drawing on an integer grid of size $(3n-2) \times (3n-2)$ with two bends per edge and one bend, respectively. The drawing can be computed in $O(n)$ time.*

Proof. We denote the outerpath by $\mathcal{Z}$ and the matching by $\mathcal{M}$. Recall that an outerpath is a biconnected outerplanar graph, whose weak dual is a path of length at least two (see Figure 6a). Let $V = \{v_1, v_2, \ldots, v_n\}$, $n \geq 4$ be the common vertex set of $\mathcal{Z}$ and $\mathcal{M}$, such that $v_1 \rightarrow v_2 \rightarrow \ldots \rightarrow v_n \rightarrow v_1$ is the outerface of $\mathcal{Z}$.

We start by triangulating all bounded faces of $\mathcal{Z}$ (without introducing new vertices). Clearly, $\mathcal{Z}$ remains outerpath but it may contain edges that also belong to $\mathcal{M}$. However, as we will shortly see this does not affect the final layout (i.e., the one derived once the extra edges introduced by the triangulation are removed). As $\mathcal{Z}$ is internally-triangulated, it contains exactly two vertices of degree two (each of which belongs to a face that corresponds to an endpoint of the dual-path). Say w.l.o.g. that $deg(v_1) = deg(v_j) = 2$ for some j with $2 < j < n$ ($j = 5$ in Figure 6a). We refer to the path $\mathcal{P}_u : v_1 \rightarrow v_2 \rightarrow \ldots \rightarrow v_j$ ($\mathcal{P}_\ell : v_{j+1} \rightarrow v_{j+2} \rightarrow \ldots \rightarrow v_n$) as the *upper* (*lower*, resp.) *path* of $\mathcal{Z}$; gray-colored in Figure 6a. Observe that, $\mathcal{P}_u \cup \mathcal{P}_\ell$ spans V. In addition, if we remove the edges of $\mathcal{P}_u \cup \mathcal{P}_\ell$ from $\mathcal{Z}$, the resulting graph is a caterpillar, say $\mathcal{C}$, that also spans V and whose spine vertices alternate between vertices of $\mathcal{P}_u$ and $\mathcal{P}_\ell$ (in Figure 6a the spine of $\mathcal{C}$ is drawn bold).

We first compute the x-coordinates of the vertices of $\mathcal{Z} \cup \mathcal{M}$. To do so, we determine the left-to-right order of the vertices of caterpillar $\mathcal{C}$, as implied by the algorithm supporting Theorem 4. Then, the x-coordinate of the i-th vertex in this order is $3i-2$, $i = 1, 2, \ldots, n$ (see Figure 6b).

In order to compute the y-coordinates of the vertices of $\mathcal{Z} \cup \mathcal{M}$, we first partition $\mathcal{M}$ into three edge-sets: $\mathcal{M}_{\ell\ell}$, $\mathcal{M}_{u\ell}$ and $\mathcal{M}_{uu}$. Edge-set $\mathcal{M}_{\ell\ell}$ ($\mathcal{M}_{uu}$) contains all edges of $\mathcal{M}$ with both end-points on path $\mathcal{P}_\ell$ ($\mathcal{P}_u$, resp.). On the other hand, $\mathcal{M}_{u\ell}$ contains all edges of $\mathcal{M}$ with one end-point on path $\mathcal{P}_u$ and the other one on $\mathcal{P}_\ell$. Since $\mathcal{P}_u \cup \mathcal{P}_\ell$ spans V, it holds that $\mathcal{M} = \mathcal{M}_{u\ell} \cup \mathcal{M}_{uu} \cup \mathcal{M}_{\ell\ell}$. In the resulting layout, edges that belong to $\mathcal{M}_{\ell\ell}$ will be drawn below the ones of $\mathcal{M}_{u\ell}$, which in turn will be drawn below the ones of $\mathcal{M}_{uu}$ (see Figure 6b). So, they will not cross with each other.

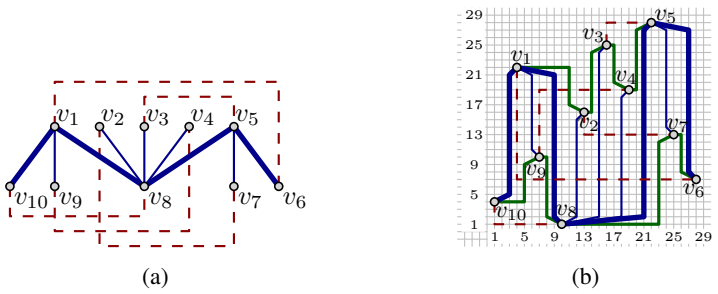

(a) (b)

Fig. 6. (a) A drawing of an outerpath and a matching. (b) A RACSIM drawing of the same outerpath and matching. In both figures the outerpath is drawn solid, the upper and the lower paths are drawn gray and dash-dotted, the spine of the spanning caterpillar is drawn bold and the matching is drawn dashed.

For $i = 1, 2, \ldots, |\mathcal{M}_{\ell\ell}|$, let $e^i_{\ell\ell} = (v_k, v_l)$ be the i-th edge of $\mathcal{M}_{\ell\ell}$, $1 \leq k, l \leq n$, and assume w.l.o.g. that $x(v_k) < x(v_l)$. Then, $y(v_k) = 6i + 4$ and $y(v_l) = 6i + 1$. In addition, edge $e^i_{\ell\ell}$ bends at point $(x(v_k), y(v_l))$. Our approach ensures that there are no crossings between edges of $\mathcal{M}_{\ell\ell}$, as they are drawn in different horizontal strips of the drawing. Similarly, we draw the edges that belong to $\mathcal{M}_{uu}$; see Figure 6b.

In order to draw the edges of $\mathcal{M}_{u\ell}$, we process all vertices that are incident to an edge of $\mathcal{M}_{u\ell}$ from left to right and we fix simultaneously the y-coordinates of both end-points of the matching edge involved. So, assume that we have processed zero or more such vertices and let $v_k \in V$ ($1 \leq k \leq n$) be the next vertex in this order that is unprocessed and its y coordinate has not been fixed. Say that $v_l \in V$ ($1 \leq l \leq n$) is the vertex that is matched with v_k in $\mathcal{M}_{u\ell}$. Then $x(v_k) < x(v_l)$. We distinguish two cases:

- If $v_k \in V(\mathcal{P}_u) \cap V(\mathcal{M}_{u\ell})$ (hence, $v_l \in V(\mathcal{P}_\ell) \cap V(\mathcal{M}_{u\ell})$), then v_k is placed three units below the bottommost vertex of $V(\mathcal{P}_u) \cap V(\mathcal{M}_{u\ell})$ that has already been placed and v_l is placed three units above the topmost vertex of $V(\mathcal{P}_\ell) \cap V(\mathcal{M}_{u\ell})$ that has already been placed[3].
- If $v_k \in V(\mathcal{P}_\ell) \cap V(\mathcal{M}_{u\ell})$ (hence, $v_l \in V(\mathcal{P}_u) \cap V(\mathcal{M}_{u\ell})$), then v_k is placed three units above the topmost vertex of $V(\mathcal{P}_\ell) \cap V(\mathcal{M}_{u\ell})$ that has already been placed and v_k is placed three units below the bottommost vertex of $V(\mathcal{P}_u) \cap V(\mathcal{M}_{u\ell})$ that has already been placed.

In both cases, edge (v_k, v_l) bends at point $(x(v_k), y(v_l))$. In addition, every edge of $\mathcal{M}_{u\ell}$ that will be drawn after (v_k, v_l) will have both of its end-points to the left of v_k and more precisely in the horizontal strip defined by the lines $y = y(v_k)$ and $y = y(v_l)$. Hence, it will not be involved in crossings with (v_k, v_l). This guarantees that $\mathcal{M}$ is drawn planar.

It remains to describe where the edges of $\mathcal{Z}$ bend. An edge (v_i, v_{i+1}) of $\mathcal{P}_u$ bends at point $(x(v_{i+1})-2, y(v_i))$ and either at point $(x(v_{i+1})-2, y(v_i)+2)$, if $y(v_i) > y(v_{i+1})$, or, at point $(x(v_{i+1}) - 2, y(v_i) - 2)$, otherwise; $i = 1, 2, \ldots, j-1$. Similarly, an edge (v_i, v_{i+1}) of $\mathcal{P}_v$ bends at point $(x(v_i) - 2, y(v_i + 1))$ and either at point $(x(v_i) - 2, y(v_i) - 2)$, if $y(v_i) > y(v_{i+1})$, or, at point $(x(v_i) - 2, y(v_i) + 2)$, otherwise; $i = j, j+1 \ldots, n-1$. An edge, say (v_k, v_l); $1 \leq k, l \leq n$, of caterpillar $\mathcal{C}$ bends either at points $(x(v_l) - 1, y(v_k) + 1)$ and $(x(v_l) - 1, y(v_l) - 1)$, if $v_k \in V(\mathcal{P}_\ell) \cap V(\mathcal{C})$ and $v_l \in V(\mathcal{P}_u) \cap V(\mathcal{C})$, or, at points $(x(v_l) - 1, y(v_k) - 1)$ and $(x(v_l) - 1, y(v_l) + 1)$, otherwise.

It is not difficult to see that both $\mathcal{P}_u$ and $\mathcal{P}_\ell$ are drawn (non-strictly) x-monotone, hence planar. Following similar arguments as in the proof of Theorem 4, we can prove that $\mathcal{C}$ is drawn planar as well. Since $\mathcal{C}$ is drawn in-between $\mathcal{P}_u$ and $\mathcal{P}_\ell$, it follows that $\mathcal{Z}$ is drawn planar, as desired. It now remains to prove that all (potential) crossings between $\mathcal{Z}$ and $\mathcal{M}$ do not involve non-rectilinear edge-segments of $\mathcal{Z}$ (as $\mathcal{M}$ consists exclusively of rectilinear segments). As all non-rectilinear edge-segments of $\mathcal{Z}$ are of unit y-length, a horizontal edge-segment of $\mathcal{M}$ cannot cross them. The same holds for vertical edge-segments of $\mathcal{M}_{uu} \cup \mathcal{M}_{\ell\ell}$, as they are drawn above and below $\mathcal{P}_u$ and $\mathcal{P}_v$, respectively. On the other hand, a vertical edge-segment of an edge of $\mathcal{M}_{u\ell}$ that

[3] Initially, we assume that the bottommost vertex of $V(\mathcal{P}_u) \cap V(\mathcal{M}_{u\ell})$ ($V(\mathcal{P}_\ell) \cap V(\mathcal{M}_{u\ell})$) that has already been placed is the bottommost vertex of $\mathcal{M}_{uu}$ (topmost vertex of $\mathcal{M}_{\ell\ell}$, resp.).

emanates, say from a vertex of $V(\mathcal{P}_\ell) \cap V(\mathcal{M}_{u\ell})$, can (potentially) cross a bundle of non-rectilinear edge-segments of $\mathcal{C}$ incident to a spine-vertex of $V(\mathcal{P}_u) \cap V(\mathcal{M}_{u\ell})$. However, in this case the spine vertex should be to the left of this vertical edge-segment and therefore by construction above it. Therefore, such crossings cannot occur, which implies all (potential) crossings between $\mathcal{Z}$ and $\mathcal{M}$ form right angles.

To complete the proof of this theorem, observe that the extra edges that we introduced while triangulating $\mathcal{Z}$ can be safely removed from the constructed layout without affecting either the crossing angles or the area of the layout. □